Rule CIC402: Records not found in CICS-maintained shared data table

Finding: The CICS interval statistics showed that a large number of records were not found in a CICS-maintained shared data table.

Impact: This finding has a MEDIUM IMPACT or HIGH IMPACT on the performance of the CICS region.

Logic flow: This is a basic finding, based on an analysis of the CICS interval statistics.

Discussion: An application can specify that a CICS VSAM key-sequenced data set (KSDS) file is to use shared data table services. When the file is opened, this specification causes CICS to copy the contents of the file into an *MVS data space*. The records can be accessed in an MVS data space significantly quicker than records read from the VSAM data set or via reads serviced by a Local Shared Resources (LSR) pool.

With shared data table support, the KSDS file is called the *source data set*. The copy of the file in memory is called the *data table*. The process of copying the records from the file to the data table is called *loading the data table*. Whenever a CICS application wishes to reference the VSAM file using normal file control commands, CICS attempts to use the representation of the file in the data table, rather than accessing the source data.

CICS supports two types of data table:

- **CICS-maintained data tables**. A CICS-maintained data table is one that CICS keeps in synchronization with their source data sets. That is, any update or delete action on a record in the data table is automatically applied to the source data set *before* being applied to the data table.

- **User-maintained data tables**. A user-maintained data table (UMT) is one that is not maintained by CICS, but is completely maintained by user code. A UMT is detached from its source data set after the table is loaded from the source data set, and changes made to the UMT are *not* reflected in the VSAM source data set.

Records are placed into a data table in one of three ways: (1) records are placed in the data table by the initial loading of the data table from the source data set, (2) records are subsequently added to the data table from the source data set, and (3) new records are written to the data table after the data table has been loaded..

- **Initial loading of the data table**. During initial loading of the data table, CICS reads the entire VSAM KSDS file and attempts to place all records in the data table. The XDTRD user exit can be used to limit the records that are placed in the data table (using screening criteria appropriate to the applications sharing the data). If all records that pass the screening criteria (if any) will not fit into the data table[1], a "table full" condition applies during the initial loading of the data table.

- **Records subsequently added from the source data set**. Records from a VSAM KSDS data set can be added to a data table after the initial loading[2] if either (1) records would not fit into the data table because the table became full during initial loading or (2) the XDTRD user exit had excluded records during initial loading but an excluded record was subsequently required by an application. If the data table had become full during initial loading, some record must be deleted from the data table for another record to be subsequently added from the source data set.

- **New records written to the data table.** New records can be added to a data table via the WRITE file control command. The XDTAD global user exit program can be used to limit the records that are placed in the CFDT as a result of a WRITE request issued to a data table (using screening criteria appropriate to the applications sharing the data). If the data table is at its maximum number of records when a WRITE is attempted, a "table full" condition applies[3].

Applications can reference the data table during initial loading. These references will produce a "record not found" condition if the references are to records outside the range of those already loaded into the data table. For both CICS-maintained data tables and user-maintained data tables, the "record not found" condition can occur during initial loading. These conditions should be small and can normally be ignored.

For user-maintained data tables[4], the "record not found" condition should not occur after loading of the data table.

For CICS-maintained data tables, small numbers of the "record not found"

---

[1] The MAXNUMRECS parameter of the DEFINE FILE command can used to limit the number of records that can be placed in the data table. The MAXNUMRECS parameter has a default maximum of NOLIMIT, which means that the entire VSAM KSDS file can be placed in the data table (subject, of course, to screening criteria applied by the XDTRD user exit).

[2] Records from the source data set cannot be added to a user-maintained data table after the initial loading of the data table. User-maintained data tables are separated from the source data set after loading.

[3] Note that this situation would not normally occur if the record had been updated, since a "read for update" would have been issued to obtain a lock on the record, and the record would simply be re-written to the table.

[4] If the "record not found" condition does occur after initial loading of a user-maintained data table, there probably is a logic or coding error with the application. These situations are analyzed in Rule CIC403.

condition after initial loading should not normally be a cause for alarm. However, a large number of "records not found" will cause unnecessary overhead and delay to applications.

- The AOR requesting the record must function ship the request to the FOR, and the record must be retrieved from the source data set. Retrieving the record from the source data set requires more overhead and more elapsed time than retrieval from the data table. Additionally, function shipping requires significantly more overhead than does cross-memory services.

  On the FOR, the source data set would be accessed to retrieve the record. Accessing the source data set for a large number of records indicates that the data table is not operating efficiently.

- If the data table size specified in the file definition is significantly less than the number of records in the source data set, only a small part of the file is loaded. Potentially, many function-shopped requests could be sent by the AOR to the FOR, not only creating unnecessary overhead but also delaying access to the records and causing increased response.

Shared data table statistics are available in MXG file CICFCR. CPExpert uses the A17DTRNF variable to assess the number of times records were not found for CICS-maintained data tables. The A17DTRNF variable contains a count of the number of times CICS attempted to read a record but was unable to satisfy the read request because the record was not in the data table; CICS was required to retrieve the record from the source data set.

The CICS statistics contain an indicator (the A17DTTYP variable) describing whether a shared data table is CICS-maintained or user-maintained. Unfortunately, this indicator is set only when the VSAM data set is closed. CPExpert does a "reverse scan" of the shared data table statistics to identify CLOSED status of a shared data table. A17DTTYP is propagated to all observations prior to CLOSE, so CPExpert can identify whether a shared data table is CICS-maintained or User-maintained. Additionally, CPExpert examines all variables that represent data set accesses that modify the data set (these are the A17DSGU, A17DSWRU, A17DSWRA, A17DSDEL, A17RMDEL, and A17DSBRU variables). If the source data set is modified, the data table must be a CICS-maintained dat table.

This approach works only if the shared data table actually was closed within the CICS statistics data available in the performance data base being analyzed by CPExpert, or if modifications were made to the data set. If the shared data table CLOSE record is not present in the data or if modifications were made to the data set, CPExpert will be unable to

determine whether the shared data table is CICS-maintained or user-maintained. From a practical matter, this should be of little import as the analysis depending on CICS-maintained versus user-maintained data tables would not be seriously effected for tables that are not closed and which have no source data set access.

CPExpert produces Rule CIC402 when the A17DTRNF value is greater than the **CICSRNF** guidance variable in USOURCE(CICGUIDE). The default value for the **CICSRNF** guidance variable is 100, indicating that CPExpert should produce Rule CIC402 when more than one hundred read requests resulted in a "record not found" condition. Note that during initial loading of the data table, any read requests outside the range of records loaded will produce a count for the "record not found" condition. As mentioned earlier, this is not normally a cause for alarm.

Suggestion: A large number of "records not found" can indicate that (1) the data table is too small for the VSAM file, or (2) the screening criteria in the XDTRD user exit does not match the record-referencing pattern of applications. If Rule CIC402 is produced regularly, you consider the following alternatives:

- Increase the number of records allowed in the shared data table. The number of records allowed in the shared data table is controlled by MAXNUMRECS parameter of the DEFINE FILE command. The DEFINE FILE command can be applied at the CEDA panel.

  Increasing the number of records in the shared data table would result in a larger table, using more virtual storage. Using more virtual storage could increase CICS paging operations, or increase overall system paging. Consequently, paging rates should be monitored if there is a large increase in the number of records allowed in the shared data table.

- If increasing the number of records allowed in the shared data table is not feasible, determine whether the XDTRD user exit has been used for the VSAM KSDS file. If the XDTRD user exit **has** been used, consider revising the screening criteria in the user exit to include a larger number of records during loading. The screening criteria might not match the records required by the application (perhaps because of different transactions or different distribution of transaction types requiring different records). Consequently, the existing screening criteria might exclude more records than are required by the applications or transaction patterns being submitted.

- Alternatively, the application might "test" for the presence of a record before adding a new record. If this is the case, the CICSRNF guidance should be changed for this specific table (Section 3 describes how to change guidance for individual shared data tables).

- You can change the CICSRNF guidance variable in USOURCE(CICGUIDE) if you believe that Rule CIC402 is produced too often.

Reference: CICS/TS for OS/390 Release 1.3 *CICS Shared Data Tables Guide*:
Section 6.1 Using the DEFINE FILE command to define data tables
Section 8.4 Interpreting data table statistics

CICS/TS for z/OS Release 2.1 *CICS Shared Data Tables Guide*:
Section 6.1 Using the DEFINE FILE command to define data tables
Section 8.4 Interpreting data table statistics

CICS/TS for z/OS Release 2.2 *CICS Shared Data Tables Guide*:
Section 6.1 Using the DEFINE FILE command to define data tables
Section 8.4 Interpreting data table statistics

